

# AUTONOMOUS WORLDS

w/

Gubsheep, Ludens, Dmstfctn, Eva Jäger,  
Alasdair Milne, Rafael Morado,  
David Huang, Bytes, Agnes Cameron,  
Guilty Gyoza, Gvn, Arb, and Ker Lee Yap

Published by OXPARC  
Edited by Guy Mackinnon-Little





## ( THE STRONGEST CRYPTO GAMING THESIS )

Gaming as a Threshold for Crypto Apps

BY GUBSHEEP

7 ♦ 24

## ( THE CASE FOR AUTONOMOUS WORLDS )

Blockchains as World Technology

BY LUDENS

15 ♦ 24

## ( LARGE LORE MODELS )

Speculative Tools for Decentralized Narrative-Building

BY DMSTFCTN, EVA JÄGER, ALASDAIR MILNE

29 ♦ 35

## ( INFINITE MODDING )

A Prehistory of Composable Game Development

BY RAFAEL MORADO

39 ♦ 45

## ( LIVELY WORLDS )

How Worlds Absorb Novelty

BY DAVID HUANG

49 ♦ 56

## ( ARCHITECTING PHYSICS )

Practical Considerations  
for the Design of Virtual Worlds

BY BYTES

61 ♦ 64



(	A CHATROOM WITH LEGS	)	@
	Indirect Communication and Conversational Worlds		@
	BY AGNES CAMERON		@
	69 ♦ 74		@
(	COMPOSABLE ENGINEERING	)	.
	Rich Digital Physics and the Evolution		@
	of Computational Worlds		
	BY GUILTYGYOZA		@
	79 ♦ 84		
(	THREE ERAS OF WORLD GENERATION	)	.
	Worldbuilding, Worlding, World-Weaving		.
	BY GVN & ARB		.
	89 ♦ 96		.
(	HOW TO FALL IN LOVE WITH A SYSTEM	)	@
	Coziness as a Catalyst for Participatory Agency		.
	BY KER LEE YAP		.
	101 ♦ 105		.
			.
			.
			.
			.
			@
			@
(	GLOSSARY	)	.
	109 ♦ 117		.
			@

AW

N1

(a)

Ⓐ

(a)

①

THE STRONGEST CRYPTO  
GAMING THESIS:  
Gaming as a Threshold for Crypto Apps

@  
@  
@  
@  
@

@  
@  
@  
@

@  
@  
@  
@  
@  
@

GUBSHEEP

Based on an essay originally written in July 2021;  
revised for Autonomous Worlds N1 in February 2023.

Crypto is not an incremental innovation. Rather, it unlocks fundamentally new ways of doing things that were not possible before. This is true across a wide set of domains: corporations reimagined as decentralized autonomous organizations (DAOs), traditional financial vehicles reimagined as permissionless DeFi protocols, legal agreements reimagined as smart contracts, and other constructs that we don't yet have the language to describe.

The most interesting and important applications of crypto will lean heavily into the new affordances provided by blockchain technology, instead of just offering incremental improvements to existing applications. This is the same pattern that we've seen with previous computing platforms: smartphones, the World Wide Web, and the personal computer itself. In each of these examples, the longest-lasting and most impactful applications of the tech use the tech to do new things, not to do old things better. Think social media apps versus newspaper websites, online stores versus digital catalogs, or interactive TV versus... well, the internet.

What does this mean for crypto gaming?

## ( CRYPTO-NATIVE GAMES )

Gaming is a leading indicator on up-and-coming technologies. Games are a technically demanding yet relatively low-stakes environment to explore scalability and usability problems early in a technology platform's lifecycle: security and compliance concerns in games are less critical compared to more 'serious' commercial or financial applications, often allowing for faster iteration loops. So while the untrained eye might dismiss a platform with 'just' a bunch of games as frivolous, the discerning will recognize that this state is a valuable testing grounds. How many mobile interaction patterns and habits have their roots in early mobile games like Doodle Jump, Cut the Rope, and

## Angry Birds?

If we accept that games are a leading indicator of new technologies, and that the most interesting applications of new technologies will lean into truly new affordances rather than incremental improvements, then it follows that the bleeding edge of crypto application design in the next few years will be found in crypto-native games.

A crypto-native game is a game that maximally embraces the architectural patterns and ethos of blockchain application development:

- The source of truth for game data is the blockchain. The blockchain is not just used as an auxiliary store of data, or a ‘mirror’ of data stored in a proprietary server. All of the meaningful data is accessible on the blockchain—not just, for example, asset ownership. This allows the game to fully utilize the benefits of a programmable blockchain: a transparent data store that is *permissionlessly* interoperable.
- The game logic and rules are implemented via smart contract. For example, combat in a game, and not just ownership, is all on-chain.
- The game is developed in accordance with open ecosystem principles. The game contracts, and an accessible game client, are open-source. Third-party developers are empowered to customize or even fork their own gameplay experiences through plugins, third-party clients, interoperable smart contracts, and even total redeployment. This in turn allows game developers to harness the creative output of an entire (incentive-aligned) community.
- The game is client-agnostic. This is closely related to the three above points, in that a litmus test for whether or not a game is crypto-native is: ‘If the core-developer-provided client disappeared tomorrow, would the game still be playable?’ The answer tends to be yes if (and only if) the game data store is permissionless, if the game logic can be executed permissionlessly, and if the community can

①

interact with the core smart contracts without relying on interfaces provided by a core team.

②

③

▪ The game is interoperable with things we consider valuable. Blockchains provide a native API into the notion of value itself, and digital assets are by default interoperable with other assets that we care about. This both reflects and contributes to a sense of depth and meaning in gameplay, and contextualizes the World of the game in the ‘real’ World.

Notably, we do not consider the following ‘crypto-native games’:

④

⑤

- A traditional mobile shooter with a crypto wallet integration.
- A TCG where card assets are stored as ERC721 tokens on Ethereum, but battles happen off-chain.
- A DeFi app that is ‘gamified’ or has a playful skin. These apps are valuable and innovative in their own right, but we are specifically interested in applications that are games first and foremost, rather than primarily financial vehicles or protocols. We consider games that are essentially pure speculation on NFTs (an asset class) to be of this category.

⑥

⑦

⑧

This is not to say that only crypto-native games can be successful (commercially or artistically), while games that use blockchain in weaker ways cannot. However, crypto-native games are the most important games to look at if we want to understand the long-term implications of blockchain technology.

( NEWMECHANICS, NEWWORLDS )

⑨

⑩


Crypto-native games are powerful because they seek to unlock mechanics that are fundamentally new, rather than just making incremental improvements on existing games.

Some of these mechanics are likely to be financial: crypto-native games certainly seem well-positioned to leverage the

powerful new incentive mechanisms and ownership models unlocked by blockchains. For example, creative and diegetic uses for tokens and in-game items could align game designers, content creators, and players in ways that weren't possible before.

But other mechanics may be social, political, or even philosophical. With an active community of only a few hundred players, Dark Forest saw levels of player self-organization that would have been unheard of in communities for online games of similar size and simplicity: cryptographic information markets, automated player corporations, embedded digital realities, and more.

Games will be instrumental in helping us to understand what blockchain technology will do for the world. As microcosms of the integrated digital Worlds of the future, they act as sandboxes for experimentation and innovation. If crypto is really going to be a civilization-scale technology, then games are a crystal ball into the future: what strange beasts will bloom in the Autonomous Worlds of tomorrow?





A polychromatic fungal colony flourishing in a terrarium project forgotten in the back closet of a high school computer lab over summer break.



·  
·  
②  
·

·  
·  
·  
②

·  
·  
·  
·  
·  
②

·  
·  
·  
·  
·  
·  
②

②

②

②

②

·  
②

②

·

·

·

·

# THE CASE FOR AUTONOMOUS WORLDS: Blockchains as World Technology

·  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a

a  
a

·  
a  
a  
a

·  
·  
·  
·  
·  
a  
a  
a

·  
·  
·  
·  
·  
a  
a  
a

LUDENS

Based off an essay originally written in August 2022; revised and extended for Autonomous Worlds N1 in February 2023.

④ A World is a container for entities and coherent-enough internal rules.<sup>(1)</sup> When a system of entities and rules becomes alive, they become a World.

④ We inhabit Worlds—whether physical or conceptual. We learn how to work and behave within them. We engage in tribalism, spatial reasoning, and territorialism over Worlds that live entirely in our minds. We have a sense for the boundaries of Worlds and their rulesets.

④ Worlds exist within *books, games, social groups, and religions*. Amongst those, we can find the Worlds of Namia, Christianity, and the Commonwealth Law.

④ Worlds run on everything from letters to wikis, bedtime stories, constitutions, databases, and, most importantly, our collective human intelligence. A World running on X means that X is the reason for the persistence of the World, the reason why we continue experiencing the container as a World, as being *alive*.

④ Worlds sometimes live entirely within the minds of people, with some light physical footprint: books, computer memory, and so on. However, the physical artefacts of those Worlds are not the reason for which the World is alive. Printing a million copies of a book doesn't create a World, unless people read it, care about it, and inhabit it.

④ ( DIEGESIS )

④ To be precise when talking about Worlds, we need to define *diegesis*. Something is diegetic *if it is in the World*. And for something to be *in the World*, it needs to have respected the *introduction rule* of the World.

④ The notion of diegesis is important when defining the boundaries of Worlds. Remember, a World is a container.

④ Let's go through some examples to build up some intuition. we will use the word *entity* to describe any constituent of a World:

events, characters, rules, facts, etc.

1. The World of Harry Potter: In the World of Harry Potter, the introduction rule is very simple: if an entity is included in a story written by JK Rowling and published under the Harry Potter series, it is diegetic. Otherwise, it isn't.
2. The World of the US dollar: This World is alien to an average person, and so is its introduction rule. Its entities are *authorities, balances, debts, and values*.

The introduction rule goes as such: If an *authority* attests the existence of a balance or debt, it is diegetic. Additionally, if enough of us accept the 'dollar value' of an entity—physical or not—its corresponding *value* becomes diegetic.

3. The World of Warcraft: In World of Warcraft, the introduction rule is formalized using computer code. If the game server relays the existence of an entity to players, it is diegetic. Introduction of new entities—like 'my character is level 60' or 'our guild is the best on the server'—is dictated by the C++ code written by Blizzard engineers.

## ( DIEGETIC BOUNDARY )

Some Worlds do not have clearly defined boundaries, and certain entities can appear to be diegetic only to a subset of people.

Most Worlds have not suffered much from ambiguities and nebulous diegetic boundaries. Others, like the USD World, are so important to our lives that we decided to spend an immense amount of time and effort on their introduction rules and borders. You can think of bureaucracy and law as a form of gravity: they attract blobs of coherent entities together and define a strict boundary for what *is* and what *isn't* diegetic.

Formalising introduction rules using 'law' and 'code' has proven to be of utmost importance to the mission-critical Worlds permeating our lives. They give Worlds *harder* diegetic boundaries.

①

② *Soft* diegetic boundaries often involve authorities or social  
 ③ consensus as a form of introduction rule, whereas *hard*  
 ④ diegetic boundaries are enforced with clear transparent rules:  
 ⑤ law, code, or mathematics. Our top level physical World—the  
 ⑥ universe—has a *very* hard diegetic boundary enforced through  
 ⑦ its introduction rule: physics.

⑧ ‘Hard’ or ‘Soft’ boundaries can both be desirable features  
 ⑨ for a World. ‘Fanfiction’ is the practice of playing with soft  
 ⑩ diegetic boundaries; while commerce requires a World with  
 ⑪ hard boundaries: doubts and arguments over the validity of  
 ⑫ someone’s tender hinders trade.

⑬ ( BLOCKCHAIN )

⑭ Worlds have a tech tree: language, writing, law, and psychology  
 ⑮ are key discoveries that enabled the creation of some of the  
 ⑯ most important Worlds around.

⑰ In 2008, an email exchange introduced one of the biggest  
 ⑱ breakthroughs in World technology: Bitcoin.

⑲ Bitcoin is a blockchain: a network technology used to create  
 ⑳ canon. In the case of Bitcoin, network participants reach  
 ㉑ agreement over the ‘canonicalness’ of a set of balances.  
 ㉒ Something being canonical is equivalent to an entity being  
 ㉓ diegetic from the lens of Worlds.

㉔ Bitcoin is a World. Just like the World of the US Dollar, it’s  
 ㉕ a weird one. The entities in the World of Bitcoin are balances  
 ㉖ and addresses, and the introduction rule is defined in computer  
 ㉗ code. It goes like this: addresses have balances. Part of a  
 ㉘ balance can be ‘spent’ to transfer it to another address.  
 ㉙ And—most importantly—balances can be increased through  
 ㉚ ‘mining,’ an expensive computational process.

㉛ Blockchains are a type of substrate for Worlds. They un-  
 ㉜ ambiguously hold the set of *all* diegetic entities within their  
 ㉝ state. Additionally, they formally define an introduction rule with  
 ㉞ computer code. A World with a blockchain substrate enables its  
 ㉟ inhabitants to participate in consensus. They run a network of  
 ㊱ computers reaching agreement on each introduction of a new

diegetic entity.

There are two blockchain concepts that are important to define from the perspective of Worlds:

1. A blockchain state root: A state root is a compression of *all* entities in the World. With a state root, one can determine whether any entity is diegetic. Believing in the state root of a World is equivalent to believing in the World itself. 0x411842e02a67ab1ab6d3722949263f06bca-20c62e03a99812bcd15dce6daf26e was the state root of Ethereum—a World with a blockchain substrate—on July 21, 2022 at 07:30:10PM UTC. All entities of the World of Ethereum were taken into account in the calculation of this state root. It represents the entirety of what was and what wasn't diegetic in that World at that specific time.
2. A blockchain state transition function: Each blockchain defines a state transition function. It can be thought of as an unambiguous introduction rule. It defines how the previous state of the World—the set of diegetic entities—can introduce new diegetic entities with inputs from people and machines. In the case of Bitcoin, the state transition function defines how balances can be spent and transferred between addresses.

In a World with a blockchain substrate, the *belief* of participants in the introduction rule entails total acceptance of the entities introduced by it. 'Belief' here needs to be defined. An inhabitant of a World with a blockchain substrate believes in the introduction rule when two statements hold:

1. They or someone they trust participate in the digital 'consensus' of that corresponding blockchain. Through participation, they can independently retrieve the state root of the blockchain, which—as described above—is a compression of all diegetic entities in the World.
2. They believe the specific *consensus algorithm* of the blockchain is operating properly. Blockchains are not magic:

they create diegetic hardness, but they aren't free lunches.

Various attacks and failure modes exist for each specific blockchain implementation.

I'd like to insist that this is *not* a default property of Worlds with formalized introduction rules. For example, a flash crash at the Chicago Mercantile Exchange led to outcomes that were rejected by almost all traders as 'invalid'; even though the introduction rule of the World—an order book matching engine—had been formalized with computer code.

Belief in the proper operation of a blockchain's consensus algorithm preempts various 'what if' scenarios found in other Worlds with formalized introduction rules.

What if someone changed the introduction rule without telling us, the inhabitants of the World?

What if the introduction rule has been misinterpreted?

What if some entities were introduced while bypassing the introduction rule?

Through digital consensus, blockchains create some of the hardest diegetic boundaries around.

( AUTONOMY )

Blockchains are of course not the only type of substrate for Worlds. Remember, Worlds run on everything from tribal songs to databases.

Yet, blockchains as a World-substrate bring a qualitative increase in the autonomy of their World.

Each World ranks differently when it comes to autonomy: some Worlds have an introduction rule relying on the existence and participation of a permissioned individual to introduce new diegetic entities (eg: Harry Potter); others rely on the consensus of a group of people to interpret and enforce their introduction rule (eg: the legal system, the World of the US Dollar); and some are in need of *untampered* computers running their formalized introduction rule (eg: The Chicago Mercantile Exchange, World of Warcraft).



In the limit case of a World's autonomy, no special individual or hardware is needed to introduce new entities and maintain the diegetic boundary.

Worlds with a blockchain substrate are maximally autonomous: anybody can enforce the introduction rule, without damaging its objectivity. The disappearance or betrayal of any particular individual does not hurt the World: its diegetic boundary remains as hard as ever. Such Worlds can be nearly on par with systems like the English language, or physics itself.

Of course, autonomy is something you can only measure in retrospect. Before an actual existential threat faces the World, autonomy is often performative. Sometimes, a credible path towards autonomy is what allows Worlds to be seen as autonomous.

## ( AUTONOMOUS WORLDS )

Given 'World with a blockchain substrate' is quite a mouthful, we started referring to them as Autonomous Worlds.

I like to think of Autonomous Worlds as planets in our solar system, but *digital* instead of *physical*.

Think about Mars. Mars—with its mountains and ancient riverbeds, its complex geology, its thin atmosphere—is a World. Most of the time, you cannot observe Mars by simply looking at the sky. Yet Mars is still out there, part of our solar system. If you were to use special instruments, you would be able to gather information about Mars, and this information would be the same for another person using the same instruments.

The telescopes used to observe Mars can be built by anyone. It makes it easier for us to agree on the fact that 'yes, there is a big red sphere out there, and you didn't make it up.'

Additionally, the rocks and deserts on Mars keep existing if someone stops believing in their World. Nobody can 'unplug' Mars.

Autonomous Worlds have telescopes that anyone can build and use to reach consensus.<sup>(2)</sup>

The entities of these Autonomous Worlds remain diegetic as

long as at least one person participates in the digital consensus. The introduction rule remains objective as well as transparent, and observing the state of the World is open to anyone with the right telescope. Nobody can unplug Autonomous Worlds.

Autonomous Worlds have hard diegetic boundaries, formalized introduction rules, and no need for privileged individuals to keep the World alive.

(FROM AUTONOMOUS WORLDS TO INTEROBJECTIVE REALITIES)

Thanks to Hilmar Petursson, Sina Habiban, and  
Guy Mackinnon-Little for inspiring this section.

In addition to our shared *objective reality* (the universe and its physics) and *our private subjective reality* (our own feelings and thoughts), we experience intersubjective realities: intangible concepts shared by multiple human beings. Prime examples of intersubjective realities are religions and money. Those realities—being subjective—have subtly different interpretations across people: *love*, an intersubjective reality, is experienced in very different ways. Even if shared, it remains intangible and subjective.

An alternative framing of Autonomous Worlds is that of ‘interobjective realities.’ Through autonomy and an objective formalized introduction rule, we can reduce—or even remove—the (inter)subjectivity of those realities.

We have taken part in intersubjective realities for tens of thousands of years. Now, using the affordances of *autonomy* and *transparency* from blockchain World-substrates, we can grant some of the rigidity and objectivity of our shared physical reality to our shared intangible realities. We can take the leap from intersubjective realities to interobjective realities.

While Autonomous Worlds offer a new way to create objective and transparent realities, it’s important to recognize that they are not meant to replace intersubjective realities. In fact, the intangibility and subjectivity of these shared concepts

are precisely what makes them so valuable and cherished by humans. However, it's crucial to understand that intersubjective realities are anchored in other realities, such as the physical World and shared cultural experiences.

In her book *The Human Condition*, philosopher Hannah Arendt talks about how common sense is like the table we sit around. She writes, 'To live together in the world means essentially that a world of things is between those who have it in common, as a table is located between those who sit around it.' Arendt suggests that the shared experience of a common human World enables us to bridge our understanding of intersubjective realities.

Understood in this way, 'common sense' might better be called 'communal sense' to distinguish it from what usually comes to mind when most of us hear the phrase. To know is to share a World, and only by sharing a common human World with others who look at it from different perspectives can we see reality in the round and develop a collaboratively authored communal sense: a canon, of sorts.

As we move away from a human World governed by physics as the underlying medium of society, our communal sense is threatened by the lack of objective shared realities. Modern digital realities are increasingly flimsy, ad-driven, AI-generated, backdoored, black-boxed, and optimized for isolation and passive consumption. It's becoming increasingly difficult to build the table, to create a communal sense with others, as the substrate of these realities keeps slipping under our feet.

Autonomous Worlds—by definition—don't slip under our feet: their Digital Physics are open and transparent, and no permissioned actor is allowed to change them without a collective agreement from their inhabitants. As our taken-for-granted intersubjective realities come under increasing threat, Autonomous Worlds can help them anchor into a more durable and stable interobjective digital reality. Autonomous Worlds can become the digital table we sit around, the container in which we can begin to fashion a new form of communal sense.

Acknowledgements: Thanks to Sina Habibian, DC Posch, Josh Stark, Saffron Huang, Rafael Morado, Hilmar Petursson, Will Robinson, Lakshman Sankar, Arthur Röing Baer, gubsheep, and Nalin Bhardwaj for feedback on the previous drafts of this essay.

Some of the ideas introduced build upon work from Ian Cheng's *Emissary's Guide to Worlding* (Metis Suns, 2018).

1. *We do not* mean worldbuilding; which is focused on creating fantasy Worlds in order to make fictional stories better and more consistent.
2. In blockchain jargon, they are called 'full nodes.'

·  
a  
a  
a  
a  
a  
a  
a  
a  
a

a  
a  
a

a

a  
a  
a  
a

a  
a

·  
·  
·  
·  
·

a  
a  
a

·  
·  
·

·  
a  
a  
a

·  
②

·  
②

②

②

·  
②

②

②

·

·

·

②

②

②

·

②

②

②

②

②

②

②

②

②

②

②

·  
②

②

·  
②

·

A formation of Roombas patrolling  
a decrepit mall.

.  
(a)  
(a)  
(a)  
(a)  
(a)  
(a)  
(a)  
(a)  
  
(a)  
  
(a)  
(a)  
  
(a)  
  
(a)  
(a)  
(a)  
  
  
(a)  
(a)  
.  
.  
.  
.  
(a)  
(a)  
(a)  
.  
.  
.  
(a)  
(a)  
(a)

②

②

②

②

②

②

②

②

②

②

②

②

②

②

②

②

②

②

②

②

②

②

②

②



LARGE LORE MODELS:  
Speculative Tools  
for Decentralized Narrative-Building

DMSTFCTN, EVA JÄGER,  
ALASDAIR MILNE

.  
@  
@  
@  
@  
@  
@  
@  
@

@

@

@

@

@

@

@

@

.

.

.

.

.

@

@

@

.

.

.

@

@

@

In the IRL world—or ‘groundworld’<sup>(1)</sup>—consensus narratives are fracturing. The Trad Web (the entrenched social networks of Web2) in particular is ineffective at delivering protocols for both ground-truthing on one hand, and healthy contestation on the other. Commentary on once verifiable events has become adversarial, with events themselves (transactions, exchanges, altercations) often difficult to verify. Here, the most fundamental record-keeping becomes centralized through ‘captured’ top-down infrastructure that resembles a surveillance architecture. The capacity to verify and contest becomes attached to hierarchy, access, and control.

While Autonomous Worlds are in their nascency, and susceptible to becoming siloed gameworlds for niche communities, they possess a wider potential as generative infrastructure for a new kind of commons—a public resource with a tangible benefit.<sup>(2)</sup> Here, we propose that decentralized narratives<sup>(3)</sup> are the essential public resource produced by Autonomous Worlds and explore the potential for Autonomous Worlds to act as a commons geared to bring this resource and its public value into the groundworld.

Autonomous Worlds can offer the ingredients and world-weaving potential to become what we term an ‘autonomous knowledge commons,’ afforded by its involuntary relationship of collective infrastructure ownership.<sup>(4)</sup> Here we take the core features of an autonomous knowledge commons to be noncentralized yet formalized access to both the means of narrative creation, linked to a permanent record of events, alongside the right to redeploy, or reseed, the results elsewhere. Building on gubsheep’s proposal that crypto-native games operate as sandboxes or ‘microcosms of the integrated digital worlds of the future,’<sup>(5)</sup> we suggest that building an autonomous knowledge commons might offer a way to export procedures or knowledge back into the groundworld in the form of narrative. With this in mind, an autonomous knowledge commons could be a space to experiment with new ways of building both consensus reality that has been lost in the groundworld, and shared narratives drawn from the emergent, permanent canon afforded

by on-chain games.

So, what would be necessary to fortify the propensity of Autonomous Worlds to become autonomous knowledge commons? Here, we will propose a new ‘modding’ superstructure: a ‘Large Lore Model’ (LLoreM) plug-in that facilitates collective and decentralized narrative-building and explicitly links to the on-chain record. The LLoreM offers an interface which maintains a link to interobjective realities<sup>(6)</sup> whilst communicating with the intersubjective groundworld. In other words, LLoreM enables a commonly built consensus reality that is nonetheless collectively generated and governed.<sup>(7)</sup> The LLoreM could be prototypical scaffolding for stewarding a shift from worlding wilderness to mission-driven commons.

## ( LORE GENERATION )

Having established the potential of Autonomous Worlds to become autonomous knowledge commons, given the right infrastructure, we look to define lore and ‘lore generation’ in concrete terms by using a procedural spoken language game, *I Went To the Shop*, as an example. From there, we will consider how to hardcode this process of lore generation into a viable tool (LLoreM) that combines the value of player lore generation with the unique affordances of on-chain games, considering the implications of such a tool in comparison to legacy Web2 lore-recording.

In the *I Went To the Shop* memory game, commonly played by kids to pass the time on road trips, we find a temporary World created through a shared narrative of a trip to the shop. The rules of the game require that players repeat the phrase: ‘Today I went to the shop and I bought...’ appending it each time with a new item purchased. The first item added must begin with the first letter of the alphabet and each subsequent item must begin with the next letter. So after three rounds the phrase might be: ‘Today I went to the shop and I bought an apple, bubblegum, and a cold beer.’ Gameplay emerges through the balance between adding items that one player can remem-

ber and items to throw off another players' memory. Players must keep a level of engagement and entertainment by listing uncommon, surreal, or offensive items. Even while the narrative produced is held in common, its texture is contingent on the particular circumstances of individual players: in-jokes, shared language (US english offers the sensible zucchini for Z, whilst in UK english you might buy a zebra), and—in the case of bored kids in a car—a desire to get a reaction from the adult driving might all play into these choices.

Here we take lore to be such a decentralized, accrued narrative. Lore generation is the general concept to describe procedures (within the World of a game or otherwise) which produce this decentralized narrative. A reified tool built for this purpose could be called a lore generator. Lore generation is important because it acts as a means of creating knowledge-claims in a suspended context, contestable later as necessary. It gamifies the production of narrative, holding it as something common and mutually constructed between players, like a jovial list of commodities bought at the shop.

But at the end of the I Went To the Shop game, the 'ledger' of ingredients has no record and is lost. Similarly, the shared experience of the players is lost, along with the emergent narrative. Further, even if some player recorded the ledger (even a trustworthy elder figure) or an on-chain version collected the list of ingredients introduced via its Digital Physics into its canon (creating a hard diegetic boundary through digital consensus), there is no mutually agreed protocol for recording and verifying the emergent narrative. In order to transform this narrative into lore then, we need a system—and interface—that is collectively agreed upon. We will now turn to a prototypical LLoreM that could be used to produce a dual ledger capable of recording and verifying a lore-claim alongside an immutable record.

## ( LARGE LORE MODELS—LLOREM )

A LLoreM is a tool we conceive of as a collective writing plug-in that attaches to a 'host' on-chain game, allowing it to

be used as what Moving Castles calls a ‘narrative engine.’<sup>(8)</sup> It is by definition a lore generator, but with additional design specifications conceptualized to respond to some of the epistemic problems discussed with regard to the groundworld. Accompanying the game’s ledger in which players’ actions are recorded through transactions on-chain, is a ‘paraledger.’ The para-ledger becomes a contestable mythology of the core on-chain gameplay events.

Writing in the para-ledger must correspond to a block on the chain, but the resultant lore is written by players and is in itself nonprocedural. The function of this design is to maximize the affordances of the blockchain’s immutability while preserving the inherent mutability of human narrative (the narrative’s contestability). The human component combines the value of subjective testimony with the incontestable actions recorded on each block. As such, the para-ledger operates as an interface through which players and observers can access and make use of the blockchain ledger, offering players an opportunity to collate or reflect upon their actions while providing observers with context beyond a list of actions. The para-ledger, of course, remains contestable: this is a core part of the design and reflects a need for intersubjective consensus in all attempts to produce collective knowledge. But any contestation at this stage, even against elder testimony, is made against the backdrop of ground-truthed canon blocks.

There are further speculative possibilities for interoperability here too: if multiple on-chain games operate on a shared blockchain ledger, and players use a singular identity to play these, the LLoreM could triangulate between the different games by locating players’ activities across the ledger, producing an ‘interdimensional’ narrative of the player as they move between Autonomous Worlds. Inversely, a lore iteration could be replayed in another game, stacking para-ledger’ lore to correspond with the blockchain and producing iterative parafiction.<sup>(9)</sup>

( COMMUNITY LOREMAXXING )

Given the stakes of narrative control we outlined at the outset, it is worth reflecting on the potential for blockchain to provide a fixed ‘ordering’ for a historical record. Narrative itself is an affordance which has been lost in the Trad Web epoch, as Lev Manovich suggests of ‘rewriting’ in Web2: ‘It is as easy to add new elements to the end of a list as it is to insert them anywhere in it. All this further contributes to the antinarrative logic of the Web. If new elements are being added over time, the result is a collection, not a story.’<sup>(10)</sup> This offers a sense of why achieving consensus reality might be hard while such a system remains dominant. By enabling packets of narrative to be anchored to the fixed-order events of an Autonomous World’s concrete, decentralized ledger, a LLoreM makes the events sense-able and transparent and thus reenables narration as a slow-burning premodern craft resistant to singular ownership by designated heroes or monarchs.

Through the specific affordances of on-chain world-weaving, the LLoreM offers an infrastructure for Autonomous Worlds to become autonomous knowledge commons. This collective lore generation system offers a procedural alternative to increasing reliance on GPT-pilled automated myth-generation. Whereas GPT models also draw upon a collective form of writing, the centralized and automated technical intervention used is much more heavily weighted, and as such fails to preserve community control over written outputs in most instances. Instead, the LLoreM seeks a way of engaging the blockchain that uses it efficiently while also aiming towards its most unique affordances; simultaneously, it is tuned to maximize the opportunities to tap into the creative potential of human collectivity: loremaxxing.

1. 'Groundworld' here refers to our IRL world that we take as ground and is collectively constituted. Alasdair Milne, *Collaborative Systems in Machine Learning Artistic Research* (PhD thesis), Serpentine Galleries R&D Platform and King's College London, forthcoming 2024.
2. Elinor Ostrom's definition of commons is a scarce resource that provides users with tangible benefits, but aren't owned by anyone. Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press, 2015. See also: Martin Zeilinger, 'Can the blockchain finally create a commons?' Spike Art Magazine, 2022.
3. We take narrative as a subjective, chronological rendering of either events in the groundworld, a fictional World, or some combination of the two.
4. *Future Art Ecosystems 3: Art x Decentralised Tech*, Serpentine Galleries, 2022. Available at: <https://www.serpentinegalleries.org/whats-on/future-art-ecosystems/>
5. gubsheep, 'The Strongest Crypto Gaming Thesis,' 2021. Available at: [https://gubsheep.mirror.xyz/nsteOfjAT-PSKH0J8IRD0j2iynmvv\\_C8i8eb483UzcTM](https://gubsheep.mirror.xyz/nsteOfjAT-PSKH0J8IRD0j2iynmvv_C8i8eb483UzcTM)
6. ludens, 'Autonomous Worlds (Part 1),' OXPARC, 2021. Available at: <https://oxparc.org/blog/autonomous-worlds>
7. Journalist and scholar Nathan Schneider has been researching cooperative models and DAOs in order to understand how crypto might offer something unique for coop tooling. Here he discusses self-governance for online communities. Nathan Schneider, 'Modpol is a Self-Governance Toolkit for Communities in Online Worlds,' Hackernoon, 2022. Available at: <https://hackernoon.com/modpol-is-a-self-governance-toolkit-for-communities-in-online-worlds>
8. Autonomous Worlds Residency Demo, 5 December 2022. See: <https://twitter.com/heylyukegibson/status/1599888301699186689>
9. Carrie Lambert-Beatty, 'Make-Believe: Parafiction and Plausibility,' October, 2009, 51-84.
10. Lev Manovich, *The Language of New Media*, The MIT Press, 2001, 220.

·  
②

·  
②  
②  
②  
·

②

·  
·  
·  
·  
·  
②  
②  
②  
·

②  
②  
②  
②  
②  
②  
②  
②  
②  
②  
②  
②  
②  
②  
②  
②  
②  
②  
②  
·

②  
②  
·  
②  
·  
·



A laboratory aquarium shimmering with programmable biomatter confounding scientists by demonstrating an entirely new form of cellular reproduction.

. AW

N1

.  
②

.  
②

②

②

.  
②

②

.

.

.

.

.

②

②

②

.

.

②

②

②

②

②

②

②

②

②

②

②

②

②

②

.  
②

.

.

INFINITE MODDING:  
A Prehistory of Composable  
Game Development

RAFAEL MORADO

·  
a  
a  
a  
a  
a  
a  
a  
a

a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a  
a

·  
·  
·  
·

a  
a  
a

·  
·  
·

a  
a  
a

In games of all kinds, there has always been the question of authorship. Gaming is a collective media and not only that, it is a performative media. A game will not unfold without the active intervention of the player, making them not an audience but a co-creator. Not all games allow for this creative performative space, but all sports and many video games do. Competition is almost always the focus of these performances (as is the case with MOBAs, shooters, and speedruns). Still, the performances of gaming can also happen in free-creating Worlds like Minecraft, whose livestreaming influencers often just explore the game for their audience. No matter the intention of the player, *playing a game is (or can be) a creative act*. While players have always been partial creators of the games they play, one way to understand the history of gaming is as a drama in which the boundaries of authorship are continually redrawn. The potentials of on-chain games herald another significant chapter in this drama, but first, let us consider how we got here.

Gaming's means of production have evolved greatly since it became a 'proper' industry over the course of the 1970s.

At the outset of the industry, one game developer would manipulate hardware directly while single-handedly creating all the content: code, graphics, music, and gameplay.<sup>(1)</sup> As time passed, more abstractions started to emerge allowing specialization within specific domains of a game. Musicians, graphic artists, and animators focused solely on executing their skillset, while engineers came to specialize in working specifically on animation tools, physics simulations, and rendering and lighting engines that allowed content creators to avoid the distraction of working with code.

Two aspects are fundamental when thinking about this evolution: ownership and complexity. As the market grew economically, big studios emerged to provide the necessary investment and retain legal ownership (both for financial liability and intellectual property) of the final product, allowing games to explode in terms of production complexity. Once these large studios started to churn out titles, it became logical to reuse large segments of code and sometimes even gameplay. Complexity

could be hidden or abstracted away, making way for standardizations. The in-house game engine thus became the path to protect the studio's investment. Most companies protected their code as IP as it could be considered a solid asset to investors. Other companies, a few years later, took a different approach and allowed not only for their code to be cloned, but for users to build on top of their game. With these 'open engines,' the stage was set for players to express themselves on a level beyond (or above) the rules and aesthetics of the game.

The modding scene that emerged from this more open approach created some of the most important titles we have today: Dota 2 and League of Legends (both iterations of the original MOBA, DoTA), Counter-Strike (built on top of Half-Life's engine), and Fortnite (an iteration of PUBG, itself a mod based on ARMA 2). Ironically, these games are now protected (owned by Epic, Riot, Valve, and Tencent), and modding them is not the simplest endeavor. Of course, not all mods are as successful as League of Legends or PUBG, but the point stands: modding enabled players to express themselves within the rules of the engine rather than only within the rules of the game. The game became to create games.

Composability on the blockchain is, philosophically, a very open type of modding. Composability, as we are defining it, is a permissionless extension of software functionality. Simple functions exist on-chain and are accessible to any user. On top of those functions,<sup>(2)</sup> more software can be created that relies on them but simplifies, amalgamates, or increases their capabilities. These new contracts are deployed freely and their use of the preexisting functions is unrestricted save for the constraints built into the initial contracts. All the mods above were created on top of existing games by people who played the game, creatively asked 'what if,' and took the effort of creating a new product. Blockchain allows this activity to happen within the live game: instead of creating a semi-clone in a separate instance, code can be deployed on top of a running game. Once games are running fully on-chain, they become fertile ground for experimentation. Anyone can deploy new features on top of a

game. These changes and additions can happen within different layers of the game, from the most fundamental (concerning the location of entities, statistics, and abilities) to the highest level (concerning the user interface, buttons, and notifications). Freedom of manipulation depends entirely on the layer and the architecture of the original game.

Let's call the bottom layer of a game the *sandbox*. At this level, the rules are there *only to enforce constraints, not goals*. Constraints are fundamental to a game because they either simplify reality so it can be simulated or they create extra friction within reality which in turn creates challenges for performative play, as is the case with a game like football. With computer simulations, we are not simplifying an existing reality, but building one from scratch. One way to think about it is that we are not only defining how much oil can be found in the virtual lands of a given game, but also how much energy burning oil provides, how to refine oil into more efficient fuels, and how to create different engines powered by these fuels (assuming those are areas the devs would expect players to delve into). Some call this set of constraints Digital Physics.

On top of the sandbox we have the *game rules*, establishing goals, rewards, and unique variants. This is where the generic constraints of the sandbox are shaped into a game, with a defined quantity of players, a starting and end condition, clear goals, tasks to be achieved, and likely a score system. The sandbox rules are indifferent to the entities they govern as they simply react in a (mostly) predictable fashion to the inputs from the user. The game rules, on the other hand, concern goals, objectives, and rewards for the users. They establish winning conditions and penalties.<sup>(3)</sup> This organization is more present in certain games (poker, sports, RPGs) and less in others (Lego or other toys). The harder the contention of these rules, the more specific skills will come out the other end. In that sense, for a skill to be measured carefully, some very tight constraints have to be put in place. Think of how enhancers are rigorously monitored in many sports.

On-chain games will necessarily abide by any rule stipulated.

As with any game played on a server, those rules cannot be escaped by the players. With on-chain games however, the developers cannot escape either. Once rules are set in motion, the player and the developer are on equal footing. When someone creates something new in the game, like a new weapon for a character, they have to adhere to the rules that are already in place. The new functionality can expand on top of the existing actions available, merge them, or synchronize them, but it cannot deny or undo them in any way. This incapable constraint is the great equalizer between developers and players, giving them both a shared potential to move the game forward. New content can now be created and deployed by anyone and the authorship of a game can evolve over time as some deployments become more popular than others. The ultimate judge of the direction a game follows is the audience, voting through play.

The constraints with which a game is deployed on-chain set the foundation for every development that follows. Given that players have the autonomy to expand the game and freely build on top of it, building solid sandboxes becomes the primary goal of those working on an on-chain game pre-deployment. A carefully designed sandbox would engage a more creative type of player, one who would be enticed to create and continuously evolve a World on top of a set of rules, not only a player focused on perfecting a narrowly defined skillset.


What, then, should these rules look like? The flexible yet simple rules of Advanced Dungeons & Dragons allow for adoption across different contexts, creating a solid base for RPGs that could evolve into uncountable versions, derivations, and campaigns. While the core mechanics were fixed, space for creation in terms of setting and narrative allowed for different games to generate enough novelty so as not to feel like they were repetitions of a previous campaign. In terms of digital games, mods like Counter-Strike or DoTA built on top of the technology of their base engine (with its collision rules, camera movements, physics, predefined classes, entity spawning and management, netcode, session creation, and client/server definitions) to

create the most popular games we have today. The creators of these popular games did not have to exhaust themselves with constructing a Digital Physics from scratch. Instead, they could experiment freely within a fixed set of constraints. As we have seen, abstraction allowed for more complex games through specialization. This standardization saved labor and made it easier to onboard developers, but more importantly, it accelerated the adoption of players who could experiment with the new game rules: as long as they had played the original game, players only had to learn what was different or new. It was both easy to introduce new game rules and easy for players to grasp them.

On-chain composability means that any game can be infinitely modded by anyone. With thoughtful design at the sandbox level, any team can introduce new functionality on top of an existing game, with the community being the final judge as to whether the functionality advances or detracts from the game. In an Autonomous World, there is no hard distinction between consumer and producer, audience and author, player and developer. Who creates the game is now truly open. Audiences can make a game their own, defying the developer's expectations. With this technology, they do not need to ask permission.



1. Most of the Atari 2600 games were developed by a single person. Some examples: David Crane (Pitfall!, Decathlon, Freeway); Carol Shaw (River Raid); Warren Robinett (Adventure).
2. We use ‘functions’ and ‘contracts’ somewhat interchangeably. It is not entirely technically accurate, but for the purpose of this article these terms refer to existing ‘smart contracts’ that are available to be executed by users. Functions are more of classic reference, where the codebase may or may not be freely accessed, while contracts are in the ‘blockchain paradigm’ and can be called by any user at any time.
3. Johan Huizinga describes this distinction between unstructured experience and the special conditions of play as the ‘magic circle,’ or ‘a state in which the player is bound by a make-believe barrier created by the game.’ Johan Huizinga, *Homo Ludens*, Beacon Press, 1955.



An agent-based model simulating the  
coordination mechanisms of foraging ants.



ΔΥΥ

(a)

(a)

(a)

(a)

(a)

(a)

(a)

•

1

•

(a)

(a)

(a)

(a)

(a)

①

①

Ⓐ

Ⓐ



①

①

a

(a)

(a)

(a)

16

10

ad

<sup>(a)</sup>

•

## LIVELY WORLDS: How Worlds Absorb Novelty

DAVID HUANG

What does it mean for something, like a world, to feel more ‘alive’ than something else? Liveliness is a vague quality, but it is nevertheless clear that when people discuss the aliveness of a thing (like a room, painting, or party), they are considering an observable characteristic. A farmers’ market on a busy weekend feels more lively to us than a virtual meeting hosted in Meta’s Horizon Workrooms.

Liveliness also seems to be irreducible to the terms of biological life, which usually involve an entity capable of self-replication and converting thermal energy into work. While it is possible to view worlds from a biological perspective, this framing can be awkward when applied to things that are not organisms. Guerrilla warfare can be alive, but not always in the same way a gorilla is alive.

The most interesting worlds are evolved, not engineered. What if we think of liveliness as being about creating space for abundant novelty and surprise? A world is most alive when it has enough open-endedness for unexpected behavior to emerge in a bottom-up, democratic manner. A lively world can be a complex system, composed of many elements that interact with each other and influence each other’s progression.

It is easy, however, to create a system that is complex, yet not maximally lively. For example, a bad codebase is complex but also frustrating enough to discourage the maintenance of itself and the World it sustains. A World that is alive has rules for how relationships between entities can be continuously created and destroyed.

## ( ON-CHAIN GAMES AND AUTONOMOUS WORLDS )

It is the nature of an Autonomous World that anyone, not only its creators, can keep it alive as long as some invariants hold.<sup>(1)</sup> Today, the most common manifestations of Autonomous Worlds are on-chain games.

A popular viewpoint to adopt in the realm of on-chain games is to magically assume that increasing agency—for example, allowing players to change the rules of the game or introduce

new entities—will automatically make a World more alive as if by some natural consequence of the blockchain medium's permissionlessness.

But in many cases, the players of a game are worse stewards of a World than its creators. Imagine an MMO game that has given players the ability to add any custom item with an associated custom behavior into the game.

Giving anyone the ability to create any type of item means everyone will make themselves gods. Players might create extremely powerful swords or impenetrable armor. But eventually, things will converge to an 'interestingness equilibrium' 'because every newly introduced object, like an invincible armor, also allows anyone to introduce a counter, like a sword that destroys an invincible armor.

This World starts lively but sinks lower on the liveliness spectrum until it is unable to surprise itself. Without formalized introduction rules, the relationships between entities in the World are left undefined and have a tendency to become meaningless.

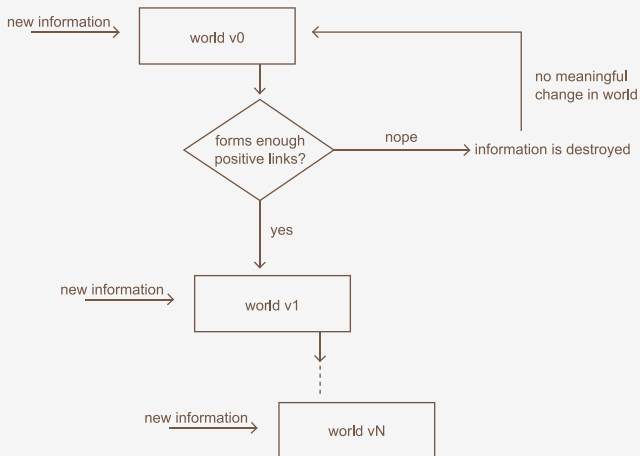
## ( FEEDBACK LOOPS IN WORLDS )

A lively World is an economy where the main currency is relationships. The most valuable relationship for encouraging liveliness is positive—where one party benefits from another without either being harmed.

Because new relationships can build on previous ones, all you need for liveliness—as continuous novelty and surprise—is a feedback loop:

1. New information is introduced to the World. This can be a discovery, like a glitch in its physics, or intentionally created, like an item for accomplishing some goal.
2. The new information disrupts relationships between entities. Existing entities might latch onto it, creating new connections and even breaking old ones.
3. If the number of positive connections the new entity receives reaches critical mass, it is 'memed into existence'

as part of the World and can interact with new information in the next iteration of the loop. Otherwise, it is destroyed by competition.



## ( COOPERATION EMERGES )

We have described a general algorithm that describes how a World changes, but have not yet defined exactly what it means for ‘relationships and connections to form or change.’ What is the specific process through which the world evaluates whether or not it can assimilate a source of novelty? How does it transition to the next version of itself?

Let us reconsider our example MMO game from before and define a more structured way for relationships to form:

- Every newly introduced item can be destroyed by anyone
- Destroying the creator also destroys the item.
- Destroyed items are expensive to rebuild.
- Each player can only introduce one item.

This creates an interface for destructive interactions—



introducing competition into the world.


Imagine if a player introduces a machine that lets you detach and reattach arbitrary limbs. If the machine unlocks new feats of heroism and spawns entire industries where bodybuilders auction off their arms and engineers can replace your legs with rockets, the World will build castles to protect the machine. But if the machine suddenly forgets how to reattach limbs, assassination markets emerge to target its creator.

How does the three-step system apply here? Players introduce new information into the World in the form of the limb machine (Step 1). The machine then interacts with existing entities in the World (Step 2). Once enough positive relationships have been formed with it, such as the emergence of a market, it ingrains itself as part of the World and is less likely to be destroyed (Step 3).

This model implies that avoiding competition is the best way entities can ensure their continued participation in this World. Suppose Player A and Player B both decide to build hyper-efficient shovels that let them harvest resources much faster. Competition emerges if this resource is scarce. Given our rules, players could try to destroy each other to escape the loop. As competition implies a constant threat, it is hard to be selfish in this World.



The other way to escape the loop is to specialize and cooperate. For example, Player A builds a device that makes crops grow larger, and works together with Player B, who keeps the hyper-efficient harvester. Cooperation can also take the form of composition, where relationships and ideas build on top of each other. This is most commonly seen in economic development, where new products or technologies are combinations of existing ones.



An item is more likely to maintain its existence when many positive links point to it. Cooperation improves each item's (and its creator's) chances of survival by increasing the number of positive links each entity receives. Consequently, the World becomes less alive when items that originate many positive relationships are destroyed. An example from nature is the extinction of a keystone species in an environment.

Each World has its own 'positive link evaluation process' that decides whether or not new information can be accepted. We can measure how well novelty incorporates itself into the World by how many positive loops—positive relationships between two entities in both directions—it creates. As long as the new information maximizes the number of positive loops it creates, there is less incentive for others to destroy it.

The drive to maximize loops means entities not only maintain positive links, but also actively seek out new relationships that

could be formed. The more positive links an entity already has, the more likely it is for newer entities to attach more links to it in the future. We can say the most lively Worlds exhibit a semilattice structure. Christopher Alexander's essay 'A City is Not a Tree' describes the semilattice as a connection pattern where all elements are deeply intertwined.<sup>(2)</sup> Alexander argues that embodying this connection-maximizing structure in our cities creates the healthiest possible communities.

## ( SYSTEMATIC WORLDING )

When tasked with the goal of creating the most realistic and interesting simulation possible, there are two approaches at opposite ends of the ladder of abstraction:

1. A symbolic approach, where every interaction and entity is defined by high-level concepts specified by humans. For example, nearly every video game.
2. A physics approach, where interactions between individual base components are represented through low-level primitives, like cellular automata. For example, falling sand simulations.

Taking the concave disposition<sup>(2)</sup>—a systematic approach to worlding<sup>(3)</sup>— will not allow us to engineer a specific outcome, but it can help us answer questions about how the different components of a world might generally work together in the future. Questions like: "Under what conditions does cooperation develop?" can be answered without zooming in on the world's implementation (physics) or zooming out on its anthropocentric biases (symbols).

You may discover that taking a systems-level view of everything makes it hard to find meaning. The heart is just a bunch of cells (muscle) that move a bunch of other cells (blood) through some more cells (arteries and veins). It does not care about the things we humans care about.

How do we understand the complexity that surrounds us?

Traditionally, cultures use storytelling and narrative to derive meaning from within a World. The heart might be a bunch of cells, but it is also the engine that moves your arms when you go to hug your family.

Meaning-making happens by identifying a slice—a journey or story—within a World and giving participants the ability to progress within that slice. The more open-ended and surprising the World is, the more opportunities there are to find the most meaningful slice for yourself. Over time, the boundaries of the slice itself (and therefore the participants' capabilities) can expand and shrink. Slices often grow into their own Worlds that bump into each other as they move through even larger meta-Worlds. The most lively Worlds have plenty of room for many journeys.

The creator of a World should view it not as a bag of systems, but as a rich medium that supports curated interfaces for meaning. As more people attempt to make Autonomous Worlds a part of our lives, we have the opportunity to elevate them from being mere containers for MMOs and drive them towards embodying models for Worlds worth living in.

1. Ludens, 'Autonomous Worlds (Part 1),' OXPARC, 2022. Available at: <https://Oxparc.org/blog/autonomous-worlds>
2. Vitalik Buterin, 'Convex and Concave Dispositions', Vitalik Buterin's website, 2020. Available at: <https://vitalik.ca/general/2020/11/08/concave.html>
3. Ian Cheng, 'Worlding Raga 2 - What is a World?', Ribbonfarm, 2019. Available at: <https://www.ribbonfarm.com/2019/03/05/worlding-raga-2-what-is-a-world/>

·  
③  
③  
③  
③  
③  
③

③  
③  
③  
③  
③  
③

③  
③  
③  
③  
③  
③  
③  
③  
③  
③

·  
·  
·  
·  
③  
③  
③  
③

·  
·  
·  
·  
③  
③  
③

·  
a

·  
a

a

a

a

a

a

a

a

·

·

·

·

·

a

a

a

a

a

a

a

a

a

a

a

a

a

a

a

a

a

a

a

a

·

a

·

·

A playground full of children crawling  
on all fours as they pretend to be ants.

. AW

N1

.

■

.

⑩

⑩

⑩

⑩

⑩

⑩

■

⑩

.

.

.

.

.

⑩

⑩

⑩

⑩

⑩

⑩

⑩

⑩

⑩

■

⑩

⑩

■

⑩

⑩

⑩

⑩

⑩

⑩

.

.

.



ARCHITECTING PHYSICS:  
Practical Considerations  
for the Design of Virtual Worlds

BYTES

This piece builds on ideas first developed in an earlier text written by bytes, cha0sg0d, and yonada.

As creators of virtual Worlds, our goal is to create environments that are engaging and interesting for users. This means finding a balance between designing a Digital Physics that allows for complex and unexpected behaviors to emerge, and ensuring that the available infrastructure can support these behaviors. To do this, we must consider three main dimensions of Digital Physics: time, the form of its laws, and the scope within which these laws apply.

We refer to the passing of time in a virtual World as the iterative application of the World's laws upon itself. Each discrete application is an 'instant' within the World's flow of time. One way to design in-World time is to have it progress continuously alongside external time. In a virtual world implemented on a blockchain, each block corresponds to a certain number of instants passing within the World, regardless of the transactions contained in the block. This is known as 'synced' time. This approach can make the World more interesting for users as it allows them to see the consequences of their actions unfold in real time. Additionally, it leads to more time passing within the World, with the World continuously updating, which in turn facilitates the emergence of interesting behaviors.

However, there are drawbacks to this approach. A larger expanse of time generally requires more computational resources, which can quickly surpass the capacity of a chain or server. It can also be difficult to implement this system on a regular blockchain, as all on-chain changes must be initiated by transactions from external users.

The alternative to synced time is unsynced time. Under this scheme, the passage of time in the World does not necessarily advance when external time does. Instead, time moves forward in response to certain events, usually user actions. Traditional board games that don't involve a timer would fall into a similar category. Unsynced time is easier to implement on-chain, as it fits the model that blockchains were designed to support.

However, it also sacrifices some of the features that could make the World more interesting.

Worldbuilders must also decide whether the mathematical laws governing their virtual World will follow open- or closed-form expressions. Closed-form expressions have a fixed number of operations. With open-form (or recursive) expressions, however, the number of operations grows depending on a given variable. Under open-form expressions, the future state of the World can only be calculated by repeatedly applying the laws of the World to a known state. Complex, live environments, such as Dwarf Fortress, usually fall into this category. Closed-form expressions, on the other hand, allow for any future state in a constant amount of time to be calculated from a past state and the time that has elapsed between them (assuming no future user actions alter the state), like pieces falling down in a game of Tetris.

Open-form expressions can make a virtual World more interesting because, like the real World, they are finitely predictable. It takes increasing amounts of time and computational resources to predict the future state of the World. Furthermore, unexpected macroscopic behaviors can emerge from simple microscopic interactions. In a World governed by closed-form expressions, these emergent behaviors generally occur only externally, through the actions of users (who themselves behave like open-form expressions), rather than within the physics of the World itself.

This trade-off between open- and closed-form expressions involves a similar balance as time. Closed-form expressions may decrease the potential interestingness of the World, but they also make it more computationally efficient. Closed-form expressions can be used with either synced or unsynced time. When implemented on a blockchain, they have a significant advantage over open-form expressions when time is synced. Because the cost of any length of time is constant, the World can be designed so that the on-chain state is only updated when a user sends a transaction, but it is set to the state it would have been in after the time since the last update had elapsed.

In the real World, time passes everywhere, all at once, in a potentially infinite universe (with some relativistic intricacies). In virtual Worlds, this is not necessarily the case.

First, the virtual World may be noticeably finite. The potential for interestingness tends to increase with size—more happens in a World made of two billion galaxies than in a World made of two atoms—but so does the computational cost. Both of these relationships are closely tied to the two trade-offs previously mentioned: the passage of time and the form of physics.

Second, time does not have to pass everywhere within the virtual World. The World can be divided into discrete regions where time passes differently, in order to reduce the computational burden of the World. For example, more sophisticated and expensive physics can be used in regions where there is user activity, while simpler physics can be used in areas where there is no activity. The downside of this approach is twofold: it can make the World seem inconsistent and lacking integrity, which limits the design space for the laws of the World and puts a burden on worldbuilders to avoid confusing users; and it also puts limits on how causality can travel within a World, as actions in one area cannot have consequences in a distant area if the space between them is frozen in time. The size of the regions where the physics apply is a major design consideration that will impact the resources needed for the World and the level of interestingness it can achieve.

To create an interesting and engaging virtual World, it is necessary to carefully balance computational efficiency with interestingness. This includes deciding on the type of time to use (sync'd or unsync'd) and evaluating the form of the physics laws that will govern the World. The size of the regions wherein the physics apply is another important decision. By making these choices carefully, not only can worldbuilders achieve interestingness while keeping the computational burden of the World manageable, they can also create a highly fertile creative substrate for other developers to build on top of.

•  
•  
•  
•  
•  
•  
•

•

•  
•  
•

•

•  
•

•

ⓐ

•  
•  
•

•  
•  
•

ⓐ

•  
•  
•

•  
•  
•

•  
•  
•

•



A speedrunning community centered around a glitch in a popular puzzle platformer, for now unpatched, that allows you to jump through walls at a perfectly perpendicular angle.

@

@

. AW

N1

.

■

.

■

■

■

Ⓜ

Ⓜ

Ⓜ

■

Ⓜ

.

.

.

.

.

■

■

Ⓜ

Ⓜ

Ⓜ

Ⓜ

Ⓜ

Ⓜ

Ⓜ

Ⓜ

■

■

■

■

■

■

■

Ⓜ

■

■

Ⓜ

Ⓜ

.

.



A CHATROOM WITH LEGS:  
Indirect Communication  
and Conversational Worlds

AGNES CAMERON

‘It may seem curious to the reader that we admit machines to the field of language and yet almost totally deny language to the ants.’

Norbert Wiener, *The Human Use of Human Beings*, 1950

@

@ In the summer of 2020, I remember spending hours in a warm haze half-watching my friend Dan play *Death Stranding*. Often somewhat unkindly referred to as a ‘walking simulator,’ *Death Stranding* is a game where the player navigates the World as a delivery worker, tasked both with supplying and gradually reconnecting a network of cities in a post-apocalyptic North America. Navigating the game, Dan would increasingly encounter items and signs left by other players—traces left in the fabric of the game, a rudimentary messaging system between ghosts, a one-sided conversation. There was an intimacy and generosity to this anonymous messaging—completing a delivery felt like picking up a sentence where someone left off, while signs and warnings became gifts passed through space-time.

@ Stigmergy describes a form of communication where the environment becomes part of the messaging medium, through fragments of memory that are spatially distributed. The classic stigmergic system is the ant colony: equipped with a small repertoire of pheromones and rudimentary sensory perception, ants communicate by leaving chemical traces for one another, a messaging system that both shapes, and is shaped by, the surrounding environment. In contrast to other forms of social organization, activity is highly decentralized—interagent communication is enough to maintain a functioning and resilient system without the need for top-down orders. Within the World of *Death Stranding*, the comparison is quite clear: ‘messages’ left for other players are almost entirely material, one-bit information that nonetheless develops an emergent complexity that proves indispensable for gameplay. However, all conversational environments, to some extent, construct and are constructed by the World in which they take place.

@ Games provide a rich territory for the discussion of environmental conversation systems. In *Dwarf Fortress*, for example,

@

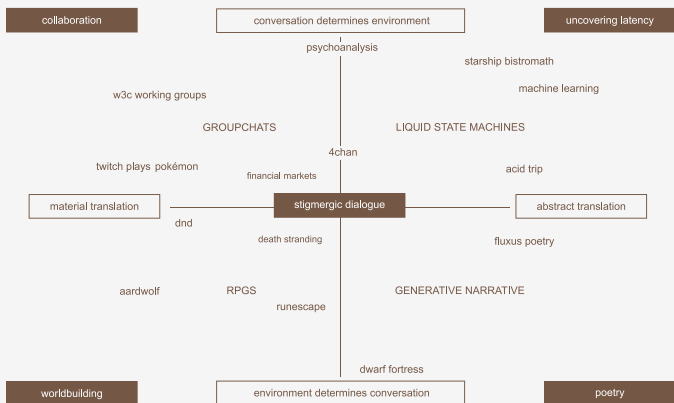
machine conversations are woven into the narrative of the game, with each character generating a long account of interactions (with other characters, plants, walls, etc.) that builds up their personality over time. As the player, however, you don't get to speak or even structure these conversations: simply setting the World in motion generates the dialogue. In many MMOs, the game exists in parallel to ongoing and obliquely related conversation—in-jokes and memes among players weave the fabric of an otherwise loosely constructed World, almost like a chatroom with legs.

What does a World look like that is explicitly rather than implicitly structured by conversation? I often come back to a scene in Douglas Adams' *The Hitchhiker's Guide to the Galaxy*, where the protagonist Arthur Dent boards a spaceship, the interior of which resembles an Italian restaurant. To his bemusement, his companion informs him that the ship is powered by Bistromathics, a field of maths that uses the complexity of interactions between patrons and staff to power a kind of random number generator, steering the ship. Would a game like this be fun to play? Perhaps not—the 'real' dynamics of the conversation, while powerful, are almost entirely hidden from the participants.

Of course, in all conversations, more is always there than is actually being said—considered in their totality, all but the most simple messaging systems contain layers of hidden meaning. The cyberneticist Gordon Pask—who coined the term 'conversation theory'—differentiates between the two dimensions of a conversation as an object-language and meta-language, the latter of which might be wholly or partially concealed from the participants.<sup>(1)</sup> The object-language is what the conversation is stated to be *about*; the meta-language constitutes a discussion of the conversation itself, which might be more or less easy for the participants to access and consciously address. I find the idea of hidden conversation dynamics within the context of feedback systems very compelling—thinking about the loops between what a particular conversation seems to be from the outside, versus what it actually *does to the World* and the people in it. In every conversation, there is the literal exchange

of information between agents, and there is the latent transformation of and feedback from their shared environment, a form of what we might term ‘stigmergic dialogue.’

While this tension is present, to some extent, in all conversation systems, I am interested here in Worlds that strike a balance between these forces of environmental and latent feedback. To me, this balancing act is akin to Fabian Fischer’s ‘Criteria for Strategy Game Design,’<sup>(2)</sup> in which he posits that an ‘interesting’ decision in a game sits somewhere between a solution and a guess. Of course, a conversation can be ‘interesting’ in terms of its subject matter—but what makes for an interesting conversation *system*? To think about this, I laid out a chart that relates what I would consider the core dynamics of stigmergic dialogue:



One is the materiality of the translation: ‘what you say is what you get’ versus systems that operate based on latent aspects of the conversation. This is similar to Pask’s object-language/meta-language dichotomy, though both the object- and meta-language are present in every conversation system to some degree. Particular to systems near the center of this axis are structures where the point of the conversation is meta-con-

versation—psychoanalysis is an obvious example, where unconscious dynamics are constantly addressed, but 4chan’s chaotic self-referentiality also follows this structure.

The other axis varies the extent to which the ‘important’ factors in an environment (note that which factors are considered important will always be subjective) are controlled by conversation versus the extent to which the environment dictates the conversation. Systems toward the center of this axis tend to have quite a rich environmental feedback system, ranging from the collaborative (DnD) to the chaotic (acid trip).

I’ve tried to characterize the ‘vibe’ of each quadrant—collaboration, worldbuilding, poetry, uncovering latency—and would argue that an interesting Autonomous World contains elements of each. What I would call ‘stigmergic dialogue’ sits at the center of these relationships: the idea of feedback between an environmental system and a conversational one, and the presence of a latent control language that participants can nonetheless participate in and perturb.

What this might mean for the design of conversation systems is an interesting question. It’s remarkable how few multiplayer games structure conversation more intentionally than a formless spatial ‘chatroom’ between players, making those that do feel quite remarkable. In *Death Stranding*, for example, a strongly stigmergic effect was achieved by heavily limiting what players could ‘say’ to one another, making the conversation entirely unidirectional and environmentally mediated. Collaborative systems like Twitch Plays Pokémon<sup>(9)</sup> and financial markets also have strictly defined rules about how communication happens. 4chan, while ostensibly totally *unconstrained* because of its anonymity and speed, actually experiences constraint for the same reasons—memes ride on waves of circular memory amongst participants; a critical mass of people ‘in the know’ is required to maintain collective energy.

In part, one could put the vibrancy of these systems down simply to ‘emergence’ (where participants in a system self-organize into a more complex structure), but a key characteristic of emergent systems is the existence of strong constraints that

permit a decentralized communication system to form. Perhaps  
 it's more interesting to argue that, by structuring a dialogue  
 system that's necessarily embedded within and constrained  
 by its environment, Worlds themselves become an emergent  
 phenomenon of communication.

1. Gordon Pask, *An Approach to Cybernetics*, Hutchinson, 1961.
2. Fabian Fischer, 'Criteria for Strategy Game Design,' Game Developer, 2014. Available at: <https://www.gamedeveloper.com/design/criteria-for-strategy-game-design>
3. Created by an anonymous Australian programmer in 2014, Twitch Plays Pokémon is a social experiment in which Twitch-users control a character by writing game-inputs directly into the chat.







A castle brimming with medieval live-action roleplayers who have gathered there to attend an imaginary feast.

AW

N1

•



•



■

(a)

■

(a)

■

■

1

1

1



(a)



■

■

1

1

■

1

1

1

1

1

1

1

1



(a)

(a)

1

# COMPOSABLE ENGINEERING: Rich Digital Physics and the Evolution of Computational Worlds

GUILTYGYOZA

## ( RICH DIGITAL PHYSICS )

Each World contains a system of fundamental laws that govern everything within it. These laws are the World's physics. Note that the term 'physics' only denotes a system of fundamental laws; it does not necessarily involve any of the familiar laws of physics from our atomic reality. In our atomic reality, if two bodies exert forces on each other, these forces have the same magnitude but opposite directions. In the World of chess, the queen may move any number of vacant squares horizontally, vertically, or diagonally.

The laws are invariants. Outside of the event of a World being upgraded, invariants stay constant and immutable. The laws form the stable structure upon which more structures at higher levels can be formed. Structures bear predictable patterns.

For Worlds that are engaging to human participants, a certain degree of predictability is required. This is because the human

mind works like a pattern-matching machine. For instance,

predictable patterns in causal effects help human participants plan ahead and make deliberate decisions. Lacking sufficient structure could lead to insufficient predictability, which leads to frustration and hurts engagement. For instance, trying to play badminton in a courtyard where wind blows chaotically from different directions would be frustrating.

We may refer to the system of fundamental laws of a computational World, or a World that is made live and interactive via computation, as its Digital Physics.

Taking the Pokémon video game worlds as an example, the Pokémon type system describes a subset of the Digital Physics of those computational Worlds.

Taking the Age of Empires video game worlds as an example, the counter system describes the relative effectiveness between military unit types. The counter system constitutes a subset of the Digital Physics of the Age of Empires computational Worlds.

The greater the number of rules structuring a World, the more interesting interactions between these rules can emerge. The

more elaborate and voluminous the body of Digital Physics becomes, the more complex processes, artifacts, and occurrences can take shape on top of it. A richer Digital Physics provides an affordance for a richer computational World.

## ( ENGINEERING AS CREATING IN THE WORLD )

With physics comes engineering: the practice of deploying the understanding of a World's physics to manipulate objects into novel and valuable configurations. Just as any design process is constrained by certain rules which concretely structure and hence enable the design itself, an engineer is both constrained and enabled by physical laws to manipulate the substances of a World into things of value. With Digital Physics, what can be engineered in computational Worlds? What *could* be?

Take the Pokémon video games as an example. Given the type system, players can engineer teams of Pokémon that are optimized to battle against particular combinations of types on an opponent's team.

Take the Age of Empires video games as an example. Given the counter system, players can engineer armies of mixed unit types that are either optimized against particular unit types on an opponent's team or optimized for complementing particular unit types from the same teams.

Players can engineer new creations in the world within the enforced boundaries of its Digital Physics.

However, for the two examples above, digging one level deeper would yield an insuperable wall: players can not engineer the individual Pokémon, nor can they engineer the military unit types. There is no physics in those computational Worlds that supports such engineering activities. New Pokémon and new military units are not engineered within the Worlds, but *introduced* into the Worlds across their diegetic boundaries by the corporate developers of those Worlds—the gods of those Worlds. This means that the superset of Pokémon is itself part of the Digital Physics of the Pokémon computational Worlds and that the superset of military units is itself part of

the Digital Physics of the Age of Empires computational Worlds.

This setup makes it difficult for these Worlds to sustain their drama, because (1) the drama of a World is partly dependent on the enumeration of objects that exist in it, and (2) this setup requires the god of the World to continue injecting new objects to sustain drama. When the enumeration of objects stays in stasis, however composable those objects might be, their combinatorics tends towards saturation. Meta—the dominant strategy to excel in the World—takes shape and becomes ossified.

Resource and power distribution among human participants tends toward stasis too. All of these effects suppress drama.

In our atomic reality, new things continually come into existence through natural evolution or by way of human discovery and invention, disrupting civilizations and societal norms, causing drama. An adaptive mutation in a virus causes global supply chains to collapse. The invention of the printing press gives rise to imaginary communities among strangers and thus the nation

state. If what exists within a World is determined by a single

corporation, the World is bottlenecked by that corporation's lifespan as well as its ability and willingness to ship—the World has reduced autonomy.

For an Autonomous World that wants sustained attention from its human participants, it needs sustained drama. For computational Worlds *off* the blockchain, the Pokémon, military units, usable equipment, consumables, vehicles, castable spells, and everything in the tech trees and skill trees are most commonly defined solely by their singular gods. All of these elements are commonly referred to as *features of a World*. For Autonomous Worlds with rich Digital Physics, they could be known as *inventions in the World*—invented from within by the World's inhabitants rather than introduced from without by its gods, keeping the World autonomous. The affordances of blockchain to enable rich Digital Physics may not be technical but cultural

and philosophical—the desire for computational Worlds that sustain themselves infinitely longer than centrally driven ones, and the rare opportunity to reinvent design approaches and business models toward sustainable worlding.

## ( COMPOSABLE ENGINEERING )

The tower of human knowledge is created by knowledge composition: the recombination of existing pieces of knowledge to unlock new epistemic and pragmatic possibilities. For example, by composing the knowledge of building a telescope and the knowledge of precise plotting through a mechanical apparatus, Galileo produced the knowledge of celestial bodies moving in ways inconsistent with what the Church asserted. This knowledge brought long-term repercussions, shaping a foundational component for nearly all physical sciences since. Human progress slows down when knowledge composition is hindered.

Composable engineering is hereby defined as the affordance of a World to allow for recursive composition of engineered artifacts with no limit on recursion depth. To give an example, the engineering of a Pokémon team produces an object with a recursion depth of zero—teams are not composable. A team is assembled to participate in battles with other individual teams; no superstructure can be built on top of a team within the confines of the Pokémon computational World. Making the system recursively composable could mean that multiple teams can be composed into a pool along with a team selection strategy, which takes an opponent team as input and returns a team from the pool that is optimally effective against the opponent team. We may call this composition of teams and selection strategy a *battle group*. To recurse one more level, imagine multiple players, each controlling one battle group, to form a regiment that battles against another regiment. In a regiment-level battle, each battle group is like a chess piece that moves as an atomic unit on a map. Special rules might govern how regiment-level resources are shared among the battle groups across the map in response to variables like morale or supply. Notice that as we recurse, game mechanics could change; game mechanics across different recursion depths could also be interdependent.

Composable engineered artifacts in Autonomous Worlds would allow for invention compounding, enabling the same process of knowledge composition that drives human history within

our atomic reality to drive the evolution of our computational Worlds. Composable engineering would also allow for knowledge encapsulation, which means, ‘I don’t need to understand every detail of your invention to involve it in my inventing process.’ Knowledge encapsulation is in some ways equivalent to the principle of separation of concerns in software development. By enabling the separation of concerns, big engineering tasks can be imagined and accomplished via chaining together small engineering tasks. Having different tasks requiring different skill sets and types of resources naturally encourages labor specialization. With labor specialization, Worlds become much more inclusive than they otherwise might be—inhabitants of different backgrounds, skill sets, and interests can all find their places in a World as creators and contributors.

This allows for diverse entries into the World, meaning more drama, more life, in the World.

As a parting thought, by involving certain cryptographic primitives in the tech stack underlying our Autonomous Worlds, information asymmetry could be introduced across the boundary of composition: ‘Not only do I not need to understand every detail of your invention, I can not possibly peak into your invention. Yet by certain quantitative measurements that are publicly available, I have confidence in the utility of your invention, hence I would trade with you to involve your invention in my inventing process.’ This asymmetry protects intellectual property rights by giving inventors an option to conceal the details of their invention and prevent free-riding forks without rendering their invention unusable.

Credits to interlocutors who contributed to my thoughts:  
Ox113d, t11s, ludens, Peteris Erins, Alan Luo







A collection of seedlings swapping nutrients in the loamy soil of a growth chamber onboard a space station orbiting Earth.

·  
·  
■  
·  
■  
■  
■  
②

·  
·  
·  
·  
■  
■  
■  
②

②

·  
·  
·  
·  
·  
·  
·  
·  
·  
·  
·  
②

②

# THREE ERAS OF WORLD GENERATION: Worldbuilding, Worlding, World-Weaving

GVN, ARB

‘Some anthropologists believe that interactive stories were invented alongside the campfire, when the elders of the tribe would incorporate audience suggestions and reactions into their performance. If this is true, then the idea of a singular author dictating the narrative is a more modern invention.’<sup>(1)</sup>

The author has been sitting alone in his chair for quite some time now, comfortably narrating the epics of great Heroes, Kings, Warriors. Of Achilles and Paris, Kenobis and Palpatines, Potters and Voldemorts, Disneys and Bezoses. Before we knew it, our ‘stories had all been pressed into service in the tale of the Hero.’<sup>(2)</sup> But it was never really *our* story, it’s always been *his*. In it, we, the audience, are to grow from the story, but never make the story grow.

We are entering an age where the audience, again, can transform the underlying structures and narrative logics of the Worlds they inhabit. Real-time rendering engines have offered a first glance into these Worlds; on-chain autonomy will slam their gates open.

To help dissipate the thick fog of war blanketing this new era, we have begun to map its fruitful yet labyrinthine lands, and compiled an archaeological record of the times preceding it.

## ( THE ERA OF WORLDBUILDING )

‘Disney (and it is no accident that his films are drawn) is a complete return to a world of complete freedom (not accidentally fictitious), freed from the necessity of another primal extinction... A fictitious world. A world of lines and colours which subjugates and alters itself to your [the animator’s] command. You tell a mountain: move, and it moves. You tell an octopus: be an elephant, and the octopus becomes an elephant. You ask the sun to stop, and it stops.’<sup>(3)</sup>

The Era of Worldbuilding is characterized by the prerendered and steadily progressing arrow of the Hero’s Journey: a circular narrative starting from the beginning and, intuitively, finishing

at the end. It is a centralized, top-down model of creation, one in which the author acts as a supreme gatekeeper of what can and cannot enter the World's walls. The author is a genius sculpting the World out of time—to use Andrei Tarkovsky's analogy<sup>(4)</sup>—the only individual with agency and control over his creation. The World itself has no autonomy to change or evolve. It remains immutable, preserved through aeons in an amber shell.

These are the immortal Worlds of Mickey Mouse, Harry Potter, and Lara Croft, hosting audiences as temporary inhabitants—or rather *tourists*—eager to consume their epic narratives.

Worldbuilding is an activity streamlined for success and capitalization at the expense of the inhabitants' agency, which is reduced to the mere passive consumption of the World's events. While this consumption might induce some emotional assimilation and produce byproducts (think of 'apocryphal' fanfiction) these will never be accepted as canon to be reintegrated into the World itself.

## ( THE ERA OF WORLDING )

'[Worlding is] the art of devising a World: by choosing its dysfunctional present, maintaining its habitable past, aiming at its transformative future, and ultimately, letting it outlive your authorial control.'<sup>(5)</sup>

We can trace the emergence of the Era of Worlding around the 1960s, when artists began to notice the potential for narratives to shift from 'the interiority of the individual [...] into a communitarian field.'<sup>(6)</sup> Characteristic of this era is the explosive, outwards motion of lightning-bolt<sup>(7)</sup> narratives—multiply authored, horizontally organized, and endlessly mutating compositions that enable a leveling, even democratizing creative process aimed at including and activating the audience. The narrative still has a beginning, but one that can now fork into myriad possible endings. With great prescience, Umberto Eco described the act

of worlding in his essay *The Open Work*<sup>(8)</sup>, where he draws on information theory to describe open artworks as those containing a multiplicity of possible interpretations. Eco's *open work* emphasizes the plurality of meaning, in which the author spawns a *field of possibilities* rather than a definitive work of art.

Through this process, the World breaks free of its amber shell, dispelling the curse that has kept it frozen in time. It becomes alive, evolving, welcoming, and attentive to the wills of its inhabitants. This ambition becomes, to us, most articulated in recent MMOs such as Ultima Online, World of Warcraft, EVE Online, and the work of artist Ian Cheng.

In these Worlds, the narrative space is opened into a communitarian field of interactions. This shift seemingly lives up to the promise of a potential return to collective authorship. And yet, the underlying logics which define the interactions, the ruleset, and virtual economies remain set in stone according to the will of their authors (now replaced with a new organizational form, the so-called *game studio*).

We argue that interaction by itself does not fulfill the promises of participatory agency if the underlying rules of a World remain unchangeable to its inhabitants. Instead of interacting with lively Worlds with real stakes and sustainable economies, players are placed in a confined sandbox, reacting to pre-established and inalterable inputs. This mirrors current critiques of and calls to democratize web platforms which, under the narrative of individual sovereignty and decentralized production of content, also prohibit the influence of underlying (often economically extractive) rulesets. A similar frustration led Vitalik Buterin to create Ethereum: a group of game developers changing the World ruleset of World of Warcraft, and the resulting extreme feeling of inadequacy that came from it.<sup>(9)</sup>

( THE ERA OF WORLD-WEAVING )

‘[...] Against a frequent misunderstanding of the notion of ‘creation’, cosmogony (world-generation) isn’t an event that



took place once and for all at the earliest point in time. The activity of worlding is repeated at every instant, in the same way that Ash'arite theologians described the world as the fragile outcome of God's continuous and arbitrary re-creation. Even the most seemingly solid and undisputable thing in the world remains vulnerable to be eradicated at any point by the twisting and turning of a subject's own metaphysical narration.<sup>9(10)</sup>

This is how the philosopher Federico Campagna describes a slightly different act of worlding, one generated through the lens of the prophet rather than the one of the author. 'A prophet,' Campagna argues, 'is not an author: it is a position towards worlding. It is a certain metaphysics, filtered through lived existence and projected as a narrative atmosphere. [...] A prophet is a place where prophecy can make itself manifest.' Blockchain World has many of these figures. Think of Satoshi: 'rendered larger than life by his anonymity, he appeared briefly as a Genesis figure and likely died a few years later. He triggered a slow revolution peopled by relatively ordinary types who bicker on Twitter and Reddit, and lack an epic center like Mecca, Silicon Valley, or Washington, DC.'<sup>11</sup> Here, epics are replaced by lore, in the form of reins that more or less gently guide the formation of a World. The World itself emerges when a set of rules laid down by the prophet are acted upon, interpreted, and reimaged by its inhabitants. The narration of these Worlds is a chaotic and decentralized one, mostly archived in wikis, discords, and game mods. The origin myth, a beginning, is still there, but instead of one or multiple ends, we might find entropy.<sup>(12)</sup>

We believe the outcome of this world-weaving to be fitting for Autonomous Worlds, a concept introduced by ludens.<sup>(13)</sup> Autonomous Worlds are 'Worlds with a clear, unalterable canon, formalized introduction rules, and no need for privileged individuals to keep it alive.' Autonomous Worlds, to us, allow a reinterpretation of the roles of Author and Audience, two roles that begin to merge into one position.

In the Worldbuilding and Worlding eras, when the inhabitants of a World misbehaved and tried to smuggle homebrew narratives into the World, the creators policed them to reestablish order.<sup>(14)</sup> In the World-Weaving era, homebrew narrative is

Ⓐ treated not as a threat to a World's order, but as the source of its renewal and continued vitality. Rather than working with scripted storylines that rigidly define a World's trajectory and harshly divide the boundaries between creators and enjoyers, we should create lore-pills and origin myths—what Ludens defines as Digital Physics—that can be assembled by the World's inhabitants and that will evolve through the tension between structure and agency.

These tensions (often used interchangeably with the concept of freedom) are a dominant theme at the center of political philosophy: from Hobbes' argument of the necessity of social order which restricts the agency of some to liberate others, to Rousseau's critique of the limits placed on agency by civil society, to Marx's arguments that social and economic structures such as class are fundamentally limiting.

Ⓐ How we design Autonomous Worlds, and according to which relation of structure and agency, ultimately becomes a political question. As the monolithic heroes and myths of the previous eras fall, so fall the systems and ideologies they were vehicles for, creating an urgency for reimagining a new politics that promotes a practice of designing Worlds where the resulting relation between structure and agency enables the value that the World produces to be captured by those whose fate is bound up with it.

Ⓐ

1. Tanya X. Short and Tarn Adams, *Procedural Storytelling in Game Design*, A K Peters/CRC Press, 2019.
2. Ursula K. Le Guin, *The Carrier Bag Theory of Fiction*, Ignota Books, 2020.
3. Sergei Eisenstein, *On Disney*, Seagull Books, 2017.
4. '[T]he film-maker from a 'lump of Time' made up of an enormous, solid cluster of living facts, cuts off and discards whatever he does not need, leaving only what is to be an element of the finished film, what will prove to be integral to the cinematic image.' Andrei Tarkovsky, *Sculpting in Time*, University of Texas Press, 1989.
5. Ian Cheng, 'Worlding Raga,' Ribbonfarm, 2019. Available at: <https://www.ribbonfarm.com/series/worlding-raga/>
6. This is how the Italian critic, curator, and historian Umbro Apollonio recognized the essential characteristic of Arte Programmata in the 1962 exhibition 'Arte programmata: Arte cinetica, opere moltiplicate, opera aperta' (Programmed art: kinetic art, multiple works, open work).
7. Jeff Gomez, 'The Hero's Journey is no Longer Serving Us,' 2017. Available at: <https://www.youtube.com/watch?v=KBBmOwe4sAU>
8. Umberto Eco, *The Open Work*, Harvard University Press, 1989.
9. Jody Macgregor, 'The creator of Ethereum got into crypto because Blizzard nerfed his character,' PC Gamer, 2021. Available at: <https://www.pcgamer.com/the-creator-of-ethereum-got-into-crypto-because-blizzard-nerfed-his-character/>
10. Federico Campagna, *Prophetic Culture*, Bloomsbury Academic, 2021.
11. Venkatesh Rao, 'Epics vs. Lore,' Ribbonfarm Studio, 2022. Available at: <https://studio.ribbonfarm.com/p/epics-vs-lore>
12. Kurt Vonnegut jokingly uses this term in his 2004 talk 'The Shape of Stories.' Available at: [https://www.youtube.com/watch?v=4\\_RUgnC1lm8](https://www.youtube.com/watch?v=4_RUgnC1lm8)
13. ludens, 'Autonomous Worlds (Part 1),' OxPARC, 2021.

Available at: <https://Oxparc.org/blog/autonomous-worlds>

14. An example of this policy is the Ultima Online ecology.

See: <https://youtu.be/KFNxJVTJleE>

@

@

@

@





A vast metafictional game in which players swap, trade, and steal a make-believe currency that has spontaneously emerged on a popular social media app.

ΔΥΥ

N1

(a)

(a)

①

①



HOW TO FALL IN LOVE  
WITH A SYSTEM:  
Coziness as a Catalyst  
for Participatory Agency

KER LEE YAP

You're finally here. Come in! The warm cafe is bustling with clusters of artists, developers, and designers in boisterous conversation. You hear people sharing ideas about their passions in gaming and worldbuilding, putting up discussion topics on a wall of sticky notes, and inviting both collaborators and oppositional opinions. Even after the symposium patrons have trickled into the autumn night, the electricity around Autonomous Worlds remains. There is some magic afoot here, and I suspect it is coziness.

As reality becomes more precarious and unpredictable, we get by through carving out cozy spaces of refuge. Nostalgia seen in the revival of retro aesthetics, pixel art, Neopets, cyberpunk, and other 1990s icons hints at the zeitgeist of returning to a more innocent and comforting time and space.

Writer Venkatesh Rao describes this trend as an age of 'domestic cozy' and homeward withdrawal.<sup>(1)</sup> The concept of coziness also appears in Maggie Appleton's 'cozy web,' which she describes as the antithesis to the manicured, publicized

'dark forest' fraught with predatory bots and trolls instead of real people.<sup>(2)</sup> The cozy web is not a treacherous forest, but a secluded garden consisting of direct messages, private servers, family photo albums, and invite-only group chats that you tend to and share with your particular niche. Coziness might be characterized by some as an escapist retreat from reality, but really it is a way of finding each other.

To get a feel for coziness' ability to empower rather than merely insulate, we can look to the recent popularity of 'wholesome' or 'cozy' games. A working group at the 2017 Project Horseshoe game design think tank produced a report on cozy games, characterizing such games as those centered on maintaining a fantasy of safety, abundance, and softness.<sup>(3)</sup> The authors posited that through strategies such as providing familiar, safe, and warmly lit environments, befriendable NPCs, collaborative leveling and gifting, and a plethora of optional side quests, developers can foster a space for trust, authenticity, and autonomy. Cozy games encourage players to linger and chat, form meaningful social bonds, and engage with the

mechanisms and systems of a game with open-ended curiosity. This is not to say that coziness must entirely define a game to be productively incorporated. More competitive, high stress games can benefit from small interludes of cozy safe zones like campfires and calm towns free from danger. The cozy qualities of these areas—access to an abundance of supplies, safety from mobs, socialization—allow players to prepare and muster courage to venture out into more challenging parts of the World. By removing immediate threats and allowing players to achieve higher order needs, cozy games might compel a player to explore what a World *might be*, rather than just try to survive what a World *already is*. By fostering ‘coziness,’ we foster participatory agency in games.

Autonomous Worlds are uniquely compatible hosts for cozy games. Autonomous Worlds promise assets and servers that run forever and with consensus, providing a sense of stability and continuity—key features of coziness—on a technical level. Decentralized Autonomous Worlds offer a place for people to put down digital roots and find refuge from the risk of their platforms being shut down or subject to changes in corporate control. In addition to technical guarantees, coziness can help cultivate the qualities of an Autonomous World that make inhabitants want to stay there long-term. Autonomous Worlds are anchored on the fact that players are simultaneously their inhabitants and creators. Anyone can contribute to the evolution and flourishing of an Autonomous World, but this affordance means nothing if the World itself doesn’t invite such participation. Coziness may serve as such an invitation. Cozy games are often associated with the pleasure of seeing emergent interactions unfold from a player’s initial inputs. Farming and gardening, for example, are hallmarks of the genre. What marks out these mechanics for the authors of the Project Horseshoe report is that they are intrinsically motivating. The joy of watching flowers bloom itself is sufficient for players to plant seeds. In the case of an Autonomous World—where players tend to the World itself rather than just a system within it—coziness can drive the care and cooperation needed for

inhabitants to design and build systems geared towards emergence and interestingness together. A cozy World is a World that you want to share with others and take care of.

We can see then how coziness can serve as a catalyst for the emergence of complex systems in Autonomous Worlds. These systems, consisting of many interconnected parts, are characterized by self-organization and the spontaneous emergence of behaviors and dynamics. The satisfaction of watching these dynamics take shape provides intrinsic motivation for players to tend to them and evolve them into new directions, which in turn provides more intrinsic motivation for further evolution. A virtuous cycle emerges where caring for the World creates a World that feels worth caring for. The comfortable

@ and secure environment created by the World's mechanics enables players to engage in sustained interactions with one another and with the World's systems. Free of external threats and pressures, players can explore non-urgent but important tasks, take more risks, and perhaps be more creative. By fostering sustained interaction and emergence, coziness can be a means of creating complex systems that are dynamic, rich, and meaningful.

As creators of new Worlds, in which longevity and liveliness are keystones, the questions often become philosophical. 'How do we get players to want to play our game?' veers into 'How do we find meaning?' Some ideas around this include introducing mechanics that promote emergent gameplay such as lively artificial intelligent NPCs, massively multiplayer experiences, decentralized servers, and blockchain-verified consensus. All of this, however, is moot if players don't feel cozy enough to actively participate and play with each other, and to keep tending to the systems and inhabitants of these worlds. The artist Ian Cheng says that the goal of his open-ended simulations is to make those who encounter them fall in love with complex systems.<sup>(4)</sup> Perhaps a similar goal can steer us in attempting to incorporate pockets of coziness into Autonomous Worlds.

1. Venkatesh Rao, 'Domestic Cozy: 1,' Ribbonfarm, 2019. Available at: <https://www.ribbonfarm.com/series/domestic-cozy/>
2. Maggie Appleton, 'The Dark Forest and the Cozy Web,' Maggie Appleton, 2020. Available at: <https://maggieapleton.com/cozy-web>
3. 'Group Report: Coziness in Games: An Exploration of Safety, Softness, and Satisfied Needs,' Project Horseshoe, 2017. Available at: <https://projecthorseshoe.com/reports/featured/ph17r3.htm>
4. Ian Cheng, *Emissary's Guide to Worlding*, Metis Suns, 2018.

a)



A cluster of sapient megaplanets  
preparing for the fabric of matter to  
fold in on itself in the last days of a  
once-flourishing intergalactic civilization.

ΔΥΥ

N1

(a)



## GLOSSARY

## AUTONOMOUS WORLDS

P. 1, 6, 12, 18, 24, 30, 36, 42, 48, 54

A world is a container for entities and a coherent-enough internal ruleset about how they behave. When a system of entities and rules comes to life, it becomes a world.

① A world is alive to the extent that its inhabitants continue to believe in its ongoing value. Artist Ian Cheng defines a world as ‘a future you can believe in: one that promises to survive its creator and continue generating drama.’

A world can take the form of a game, a community, an institution, a religion, a cosmology, and so on.

( WORLD )  
 ≈ Narrative, Story, Tale

P. 12, 36, 48

The stories told about a world by those who inhabit and interact with it.

Lore encompasses the entire scope of a world's collective narrativization both within and beyond its formalized canon.

@  
@

( LORE )  
≈ Mythology, Collective Knowledge

P. 6, 12, 18, 24, 30, 36, 42, 48, 54

The introduction of new information into the world by human inhabitants or automated entities, and the form this introduction takes. Whether or not an input is accepted by a world is determined by the nature of its diegetic boundary.

( INPUT )  
 ≈ Turn, Move

P. 6, 18, 24, 42

The boundary condition that defines what is considered part of a world and what is not. The diegetic boundary of a world can be variously soft—enforced via subjective belief or social consensus—or hard—enforced via law, code, or mathematics.

( DIEGETIC BOUNDARY )  
 ≈ magic circle, introduction rule

P. 6, 18, 24, 30, 42

A set of invariable rules that govern everything that takes place within a world. That which determines the causal effects from a world's inputs.

@

@

( DIGITAL PHYSICS )  
 ≈ Ruleset, Structure, Laws

@

P. 1, 6, 12, 24, 48

A historical log of accepted inputs  
and their causal effects according to  
a world's digital physics.

( CANON )  
≈ History, Records

P. 6, 18, 42

A world's ability to be extended by anyone that respects its internal ruleset.

( COMPOSABILITY )  
 ≈ Permissionlessness, Modularity



P. 1, 6, 12, 18, 24, 30, 36, 42, 48, 54

A world with a clear, unalterable canon, formalized digital physics, and no need for privileged individuals to keep it alive.

( AUTONOMOUS WORLD )  
 ≈ Hyperstructure, Crypto-Native game

.  
◇

.

■

■

■

■

■

■

◇

■

.

.

.

Ⓐ

Ⓐ

.

◇

■

.

.

■

■

■

■

■

■

◇

◇

■

◇

◇

■

■

■

■

.

.



The essays in this publication are the result of an open-call sent to participants of the Autonomous Worlds Symposium, hosted in London 2022.

Published by OxPARC

Editor: Guy Mackinnon-Little

Art direction: Leïth Benkhedda

Creative Producer: Arthur Röing Baer

Code (ASCII cellular automata): Son-La Pham

Illustrations: PWR studio

Typeset in FK Raster Grotesk

Printing and binding: Druckhaus Sportflieger, Berlin

ISBN: 978-1-3999-4967-5

Oxparc.org

